

Competitive Capacitated Online Recoloring

Rajmohan Rajaraman  

Northeastern University, Boston, MA, USA

Omer Wasim  

Northeastern University, Boston, MA, USA

Abstract

In this paper, we revisit the *online recoloring* problem introduced recently by Azar, Machluf, Patt-Shamir and Touitou [5] to investigate algorithmic challenges that arise while scheduling virtual machines or processes in distributed systems and cloud services. In online recoloring, there is a fixed set V of n vertices and an initial coloring $c_0 : V \rightarrow [k]$ for some $k \in \mathbb{Z}^{>0}$. Under an online sequence σ of requests where each request is an edge (u_t, v_t) , a proper vertex coloring c of the graph G_t induced by requests until time t needs to be maintained for all t ; i.e., for any $(u, v) \in G_t$, $c(u) \neq c(v)$. In the distributed systems application, a vertex corresponds to a VM, an edge corresponds to the requirement that the two endpoint VMs be on different clusters, and a coloring is an allocation of VMs to clusters. The objective is to minimize the total weight of vertices recolored for the sequence σ . In [5], the authors give competitive algorithms for two polynomially tractable cases – 2-coloring for bipartite G_t and $(\Delta + 1)$ -coloring for Δ -degree G_t – and lower bounds for the fully dynamic case where G_t can be arbitrary.

We obtain the first competitive algorithms for *capacitated* online recoloring and fully dynamic recoloring, in which there is a bound on the number or weight of vertices in each color. Our first set of results is for 2-recoloring using algorithms that are $(1 + \varepsilon)$ -resource augmented where $\varepsilon \in (0, 1)$ is an arbitrarily small constant. Our main result is an $O(\log n)$ -competitive *deterministic* algorithm for weighted bipartite graphs, which is *asymptotically optimal* in light of an $\Omega(\log n)$ lower bound that holds for an unbounded amount of augmentation. We also present an $O(n \log n)$ -competitive deterministic algorithm for *fully dynamic* recoloring, which is optimal within an $O(\log n)$ factor in light of a $\Omega(n)$ lower bound that holds for an unbounded amount of augmentation.

Our second set of results is for Δ -recoloring in an $(1 + \varepsilon)$ -overprovisioned setting where the maximum degree of G_t is bounded by $(1 - \varepsilon)\Delta$ for all t , and each color assigned to at most $(1 + \varepsilon)\frac{n}{\Delta}$ vertices, for an arbitrary $\varepsilon > 0$. Our main result is an $O(1)$ -competitive randomized algorithm for $\Delta = O(\sqrt{n/\log n})$. We also present an $O(\Delta)$ -competitive deterministic algorithm for $\Delta \leq \varepsilon n/2$. Both results are *asymptotically optimal*.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases online algorithms, competitive ratio, recoloring, resource augmentation

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.95

Related Version *Full Version*: <https://arxiv.org/pdf/2408.05370>

Acknowledgements This work was partially supported by NSF grant CCF-2335187.

1 Introduction

The challenge of efficiently allocating virtual machines (VMs) across large clusters in a cloud or data center is well documented [26, 24, 27]. Most big-data applications are inherently distributed. To minimize congestion, communication costs, and service delays incurred due to inter-VM and inter-process communication, distributed systems are demand-aware [4], requiring scheduling algorithms to efficiently adapt to network traffic. While co-locating frequently communicating VMs and processes in a single cluster (corresponding to affinity requests in cloud or VM environment) is desirable to minimize communication overhead, service providers are also required to satisfy anti-affinity requests [1, 16], which require that certain VMs must not be co-located. Large-scale deployment systems including Kubernetes and VMWare support anti-affinity rules to provide safety, performance and robustness [2, 10].



© Rajmohan Rajaraman and Omer Wasim;
licensed under Creative Commons License CC-BY 4.0
32nd Annual European Symposium on Algorithms (ESA 2024).

Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 95; pp. 95:1–95:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The online recoloring problem (also referred to as online disengagement) was introduced recently by Azar, Machluf, Patt-Shamir and Tseitou [5] to investigate algorithmic challenges that arise in scheduling of VMs in distributed systems and cloud computing. In this problem, we have a graph G on a fixed set of n vertices (corresponding to VMs) with weights (corresponding to VM sizes) and edges (corresponding to anti-affinity requests) are revealed over time. The goal is to ensure that endpoints of every edge encountered are assigned a different color (i.e. the corresponding VMs scheduled on different clusters). The graph induced by the requested edges is assumed to be k -colorable (colors corresponding to k clusters) and the objective is to minimize the total weight of vertices recolored throughout time. The cost of recoloring any vertex captures the size and migration cost of the corresponding VM.

In [5], the authors study the *uncapacitated* version where the total weight of vertices assigned a color is unconstrained. Since vertex coloring is NP-hard [17], they focus on polynomially tractable cases when the graph induced by the request sequence is either bipartite ($k = 2$) or has maximum degree at most Δ ($k = \Delta + 1$). In the static setting, both cases admit linear-time *greedy algorithms*. In the online setting, while existence of a proper k -coloring guarantees that any vertex needs to be recolored at most once, the uncertainty of future requests can force multiple recolorings. The *capacitated* online recoloring problem captures the practical scenario where cluster sizes are constrained, i.e. the total weight of vertices assigned any color is bounded. Resolving this was left as an open problem in [5].

Online Recoloring. We formally describe the online recoloring problem [5]. Given a fixed set of n vertices denoted by V , weight function $w : V \rightarrow \mathbb{R}^+$ and an initial coloring $c_0 : V \rightarrow [k]$ (where $[k] = \{1, 2, \dots, k\}$), an algorithm maintains a k -coloring $c : V \rightarrow [k]$ for V under an online sequence of requests, $\sigma = \{\sigma_i\}_{i=1}^T$ where σ_i is an edge (u_i, v_i) between $u_i, v_i \in V$. Let $\sigma_{\leq t} = \{\sigma_i\}_{i=1}^t$. For all t , c must be a proper k -coloring for G_t , the subgraph induced by $\sigma_{\leq t}$, i.e. for all (u_i, v_i) where $i \leq t$, $c(u_i) \neq c(v_i)$. At any time, any vertex v can be recolored for a cost $w(v)$. In the unweighted setting, $w(v) = 1$ for all $v \in V$. The objective is to minimize the total recoloring cost incurred for the request sequence σ .

We analyze online algorithms in the standard competitive analysis framework. Let the total cost by an optimal algorithm (resp., an online algorithm \mathcal{A}) on a request sequence σ be denoted by OPT_σ (resp. \mathcal{A}_σ). An algorithm \mathcal{A} for online recoloring is α -competitive if for any request sequence σ , $\mathcal{A}_\sigma \leq \alpha OPT_\sigma$. Note that σ can be assumed to be finite w.l.o.g. (in particular, $|\sigma| = O(n^2)$) since a repeated edge request incurs no cost.

Fully Dynamic Recoloring. In the above formulation, the graph induced by the request sequence is assumed to be k -colorable. We also consider a *fully dynamic* variant [6], which captures the case when requests are temporary, i.e. the graph induced by the request sequence is not necessarily k -colorable. For all t , it is *only* required to maintain that $c(u_t) \neq c(v_t)$ for the current request $\sigma_t = (u_t, v_t)$. Note that the request sequence can be unbounded in this model. An algorithm \mathcal{A} for fully dynamic recoloring is α -competitive if for any request sequence σ , $\mathcal{A}_\sigma \leq \alpha OPT_\sigma + \mu$ where μ is a constant independent of request sequence σ .

Capacitated Online and Fully Dynamic Recoloring. For *uncapacitated* recoloring described above, the total weight of vertices assigned a particular color is unconstrained. In *capacitated* recoloring, the total weight of vertices assigned color i is at most W_i for all i at any time t . In this paper, we focus our attention on the *uniform* capacity case, i.e. $W_i = B$, for all $i \in [k]$. In the unweighted case, B is the *number* of vertices assigned any color $i \in [k]$. For capacitated online and fully dynamic recoloring, we assume that the graph induced by σ admits a proper k -coloring while respecting capacity constraints for each $i \in [k]$.

Resource Augmentation. Our algorithms for online and fully dynamic 2-recoloring *mildly utilize* resource augmentation: we assume the total weight of vertices that can be assigned color i is $B(1 + \varepsilon) \forall i \in [k]$ for an arbitrarily small constant $0 < \varepsilon < 1$. The competitive ratio is analyzed with respect to an offline algorithm constrained to a capacity of B for all i . Nevertheless, our deterministic and randomized lower bounds hold for resource augmented algorithms—thus, establishing tightness of the upper bounds we obtain.

Overprovisioned Setting. For Δ -recoloring, we introduce and study an *overprovisioned* setting where the graph induced by σ has maximum degree $(1 - \varepsilon)\Delta$ and each color $i \in [\Delta]$ has capacity $(1 + \varepsilon)n/\Delta$. In contrast to the resource augmented framework, the competitive ratio is analyzed with respect to a offline algorithm with the same number of colors and unconstrained capacity. Remarkably, even in this pessimistic analytical framework, we obtain tight deterministic and randomized algorithms.

1.1 Related Work

Online Recoloring. Azar et al. [5] recently introduced the online (uncapacitated) recoloring problem. They give optimal deterministic and randomized algorithms for online recoloring when $k = 2$ and $k = \Delta + 1$. For the bipartite case, they give an optimal $O(\log n)$ -competitive deterministic algorithm and a lower bound of $\Omega(\log n)$, which holds for randomized algorithms. For $(\Delta + 1)$ -recoloring, they give a deterministic $O(\Delta)$ -competitive algorithm, a $O(\log \Delta)$ -competitive randomized algorithm, and matching lower bounds. They introduce fully dynamic recoloring and consider three equivalent edge insertion models and show that fully dynamic recoloring has substantially worse competitive ratios than online recoloring: for 2-coloring, they consider edge requests on a length n -odd cycle (which cannot be properly colored in the online recoloring setting) and provide a reduction from the metrical task systems (MTS) problem on a length n -odd cycle, yielding $\Omega(n)$ deterministic and $\Omega(\frac{\log n}{\log \log n})$ randomized lower bounds respectively for fully dynamic recoloring.

In subsequent work [6], Azar et al. give general lower bounds for k -coloring. They provide $\Omega(k \log(n/k))$ deterministic and $\Omega(\log k \cdot \log(n/k))$ and randomized lower bounds. They also present lower bounds for capacitated k -coloring, including an $\Omega(n)$ deterministic and an $\Omega(\log n)$ randomized lower bound, for the *tight* case when $kB = \sum_{v \in V} w(v)$. Note that all lower bounds given in [6, 5] hold even for *unweighted* instances, i.e. when $w(v) = 1, \forall v \in V$.

Dynamic Coloring. Maintaining a proper coloring in fully dynamic graphs is well studied [14, 15]. Since determining the chromatic number of a graph is NP-hard [17], most works focus on restricted classes of graphs, bounded chromatic number or maximum degree. In dynamic coloring, an edge (or vertex) is either inserted or deleted from the graph and the algorithm is required to quickly re-compute a proper coloring. The objective is to minimize the update time [8, 25], and/or the recourse [18, 7] (i.e. the number of recolored vertices after an update) in a worst-case or amortized sense. Kashyop et al. [18], Barba et al. [7] and Bosek et al. [9] study trade-offs between the number of colors, number of recolorings and update time for dynamic coloring for various classes of graphs (including bipartite, bounded degree and arboricity, and interval graphs). The crucial difference between dynamic and online recoloring is the fact that the number of recolorings in the former do not guarantee competitiveness in the latter; a dynamic algorithm incurring $O(1)$ recolorings per update only guarantees $O(T)$ recolorings where $T = |\sigma|$ in the online setting. When $OPT_\sigma \ll T$, this precludes competitiveness which is measured against OPT_σ .

Online Balanced Graph Partitioning. Online balanced graph partitioning (OBGR) introduced by Avin et al., [3] and later explored in [22, 13, 21, 23] is related to online recoloring. While the objective of recoloring is to optimize allocations of VM's under *anti-affinity* requests, OBGR deals with *affinity* requests. In OBGR, given a graph G on n vertices that are initially assigned to ℓ servers each of capacity k , for any online request (u, v) , an algorithm incurs a unit cost if u and v are in different clusters and 0 otherwise. Any vertex can be migrated for a unit cost at any time. The goal is to minimize the total migration and communication costs incurred for the request sequence. Since the static version of balanced graph partitioning is known to be NP-hard even for $\ell = 2$ (the minimum bisection problem) [12] resource augmentation and randomization have been employed to obtain competitive algorithms (some taking exponential time) for OBGR. Avin et al. [3] give an $O(k^2\ell^2)$ competitive algorithm without resource augmentation, and a $O(k \log k)$ exponential time algorithm with $(2 + \varepsilon)$ -augmentation (improved in [11] with $(1 + \varepsilon)$ -augmentation). For the *learning* variant of OBGR, Henzinger et al. [13] give optimal deterministic $O(\ell \log k)$ - and randomized $O(\log k + \log \ell)$ -competitive algorithms. Räcke et al [22] give an $O(\log n)$ -approximation algorithm with $(2 + \varepsilon)$ -augmentation for the *offline* case. The current best deterministic online algorithm for OBGR [23] obtains $O(k\ell \log k)$ -competitiveness with $(1 + \varepsilon)$ -augmentation.

1.2 Our results

We present the first competitive algorithms for *capacitated* online 2-recoloring, making progress on open problems posed by Azar et al. [5]. We also introduce the capacitated Δ -recoloring problem in an overprovisioned setting and obtain asymptotically optimal algorithms.

Capacitated 2-Recoloring. Our algorithms for capacitated 2-recoloring hold with $(1 + \varepsilon)$ -resource augmentation, i.e. the algorithms utilize a capacity of $W = (1 + \varepsilon)B$ for each cluster, while the optimal offline algorithm has a capacity of B where $B = \frac{\sum_{v \in V} w(v)}{2}$, (w.l.o.g., assume $\sum_{v \in V} w(v)$ is even) and $\varepsilon \in (0, 1)$ is a small constant.

Fully Dynamic Recoloring: Our first result is an algorithm for *fully dynamic* capacitated recoloring which works for unweighted instances.

► **Theorem 1.** *There exists a deterministic $O(n \log n)$ -competitive $(1 + \varepsilon)$ -resource augmented algorithm for fully dynamic capacitated online 2-recoloring that works for unweighted instances.*

We give a deterministic lower bound for fully dynamic 2-recoloring that holds for an unweighted instance and *unbounded* resource augmentation, establishing that our algorithm is tight within a $O(\log n)$ factor. The proof of Theorem 2 is deferred to the full version of the paper.

► **Theorem 2.** *The competitive ratio of any deterministic algorithm for capacitated fully dynamic 2-recoloring is $\Omega(n)$, for any resource augmentation $\varepsilon > 0$ and unweighted instances.*

Online Recoloring: Our main result is an asymptotically *optimal* deterministic $O(\log n)$ -competitive algorithm for capacitated online 2-recoloring which works for *weighted* instances.

► **Theorem 3.** *There exists an optimal deterministic $O(\log n)$ -competitive $(1 + \varepsilon)$ -resource augmented algorithm for capacitated online 2-recoloring that works for weighted instances.*

We complement our result giving a lower bound for the capacitated case similar to the uncapacitated case in [5], which holds for an *unbounded* amount of resource augmentation, $\varepsilon > 0$. The proof of Theorem 4 is deferred to the full version of the paper.

► **Theorem 4.** *The competitive ratio of any randomized algorithm for capacitated online 2-recoloring is $\Omega(\log n)$, which holds for an arbitrary amount of resource augmentation $\varepsilon > 0$.*

Capacitated Δ -recoloring. We introduce the $(1 + \varepsilon)$ -overprovisioned setting for capacitated Δ -recoloring, and give optimal deterministic and randomized algorithms. In this setting, there are n vertices, and Δ colors. Each color is assigned exactly $\frac{n}{\Delta}$ vertices initially. The maximum degree of the graph induced by the request sequence is bounded by $(1 - \varepsilon)\Delta$, and the total capacity per color is $(1 + \varepsilon)\frac{n}{\Delta}$. The request sequence is finite. In contrast to the resource augmented setting for which the competitiveness is analyzed with respect to an unaugmented optimal offline algorithm, we analyze competitiveness with respect to an optimal offline algorithm with the same number of colors and unrestricted capacity per color.

We first note that a lower bound of Δ holds for any deterministic algorithm in this setting. The instance is similar to the one given for uncapacitated $(\Delta + 1)$ -recoloring given in [5]. Thus, deterministic $(1 + \varepsilon)$ -overprovisioned Δ -recoloring is as hard as deterministic $(\Delta + 1)$ recoloring. The proof of Theorem 5 is deferred to the full version of the paper.

► **Theorem 5.** *The competitive ratio of any deterministic algorithm for $(1 + \varepsilon)$ -overprovisioned Δ -recoloring is $\Omega(\Delta)$.*

We give an optimal deterministic algorithm matching the above lower bound.

► **Theorem 6.** *There exists a $O(\Delta)$ -competitive deterministic algorithm for $(1 + \varepsilon)$ -overprovisioned Δ -recoloring when $\Delta \leq \varepsilon n/2$.*

Our main result for Δ -recoloring is an optimal $O(1)$ -competitive randomized algorithm.

► **Theorem 7.** *There exists a $O(1)$ -competitive randomized algorithm for $(1 + \varepsilon)$ -overprovisioned Δ -recoloring that works against an oblivious adversary when $\Delta = O\left(\sqrt{\frac{n}{\log n}}\right)$.*

1.3 Technical Overview

Algorithms for 2-Recoloring. Our algorithms maintain a set of bipartite components induced by the request sequence. We develop a subroutine **Rebalance** which is a FPTAS, to periodically assign components while respecting capacity constraints.

Our fully dynamic algorithm is a phase-based algorithm where the graph induced by requests in any phase is bipartite and admits a coloring such that the total weight of vertices assigned any color is $W = (1 + \varepsilon)B$. During a phase, our algorithm works as follows: on a request (u, v) where u is in component P_1 and v in P_2 , s.t. P_1 is the heavier of the two components, P_2 is recolored and merged to P_1 . If recoloring of P_2 leads to capacity violation, **Rebalance** is invoked. The total recoloring cost charged to a vertex per phase is bounded by separately considering when it is part of a heavy or light component. We observe that whenever recoloring of a light component leads to a call to **Rebalance**, vertices in light components of total weight at least $\Omega(W)$ must have been recolored since the last call to **Rebalance** (if any). On the other hand, when two heavy components are merged, **Rebalance** is always invoked; while a single call to **Rebalance** incurs $O(W)$ cost, the number of such calls is only $O(1)$ per phase. Combined with a charging argument which bounds the number of times any vertex can be recolored in a phase by $O(\log n)$, we obtain a $O(n \log n)$ -competitive algorithm for fully dynamic recoloring. We complement this result with a lower bound of $\Omega(n)$ which holds for an unbounded amount of augmentation, establishing that our algorithm is tight to within a $O(\log n)$ factor.

Our deterministic $O(\log n)$ -competitive algorithm uses the fully dynamic algorithm as a subroutine once a lower bound of $\Omega(W)$ on OPT_σ has been determined. Our algorithm incorporates *laziness* together with the *greedy* approach. While it is crucial to maintain a coloring that is *close* to the initial coloring for any component to be competitive, we recolor a component according to an optimal coloring only after a constant factor increase in its weight, until it is possibly merged to a heavier component. Whenever a light component P_2 is merged to a heavy component P_1 such that the weight of P_1 does not change significantly, only P_2 is recolored (if necessary). A crucial property we exploit in the analysis is the monotonicity of the distance to initial coloring with respect to the weight of any component. As a result, the cost of periodically re-coloring components to mimic an optimal coloring after weight increases can be charged to OPT_σ . An intricate charging argument allows us to bound the cost of both periodic recolorings and greedy recolorings by $OPT_\sigma O(\log n)$.

If the request sequence precludes optimal colorings without significant rebalancing (i.e. when a lower bound of $OPT_\sigma = \Omega(W)$ is determined), our algorithm transitions to the fully dynamic algorithm. The total cost of the fully dynamic algorithm for the remaining request sequence is bounded by $O(W \log n)$. Note that our fully dynamic algorithm is $O(n \log n)$ -competitive. Our analysis reveals that in the case when $OPT_\sigma = \Omega(W)$, the fully dynamic algorithm yields $O(\log n)$ competitiveness on *weighted* instances. We give a *randomized* lower bound of $O(\log n)$ for online recoloring which holds for an algorithm with *unbounded* resource augmentation, establishing that our algorithm for online 2-recoloring is *asymptotically optimal*.

Algorithms for Δ -Recoloring. We give *asymptotically optimal* deterministic and randomized algorithms for Δ -recoloring in the $(1 + \varepsilon)$ -overprovisioned setting. Both algorithms consist of a *recoloring* subroutine that colors an individual vertex and a *rebalancing* subroutine when any color reaches capacity. The implementations of these subroutines are different in the two algorithms. Similar to the approach in [5], we lower bound OPT_σ as the size of the minimum vertex cover on the graph induced by the set of monochromatic edges in σ with respect to the initial coloring. Both algorithms maintain a 2-approximate vertex cover C . If the requested edge is monochromatic, both algorithms recolor only the endpoint in C .

Our deterministic algorithm picks a feasible color assigned to the least number of vertices, ensuring that whenever all feasible colors are *full* for a vertex, the rebalancing cost of $O(n)$ can be charged to the increase in vertex degrees. To rebalance, we utilize an algorithm for equitable coloring of graphs with bounded maximum degree by Kierstead et al. [19]. We complement our results by giving a deterministic $\Omega(\Delta)$ lower bound, similar to an instance in [5].

Our $O(1)$ -competitive randomized algorithm uses a simple rebalancing procedure which sequentially colors vertices randomly among their feasible colors, where a color c is feasible for a vertex v if none of v 's neighbors are assigned c . To implement the recoloring subroutine for v , a random feasible color for v is picked. While our randomized algorithm is simple, we face multiple technical challenges towards obtaining competitiveness since straightforward concentration inequalities cannot be applied, because recolorings (and rebalancings) involve *dependent* events. For instance, our analysis of rebalancing utilizes a martingale to derive a concentration bound on the load of any color. We also invoke stochastic domination arguments to upper bound the load on any color due to recoloring, and lower bound the number of recolorings performed before any color is overloaded.

1.4 Open Problems

While we obtain tight bounds for several variants of capacitated recoloring, our work leaves several open problems. Obtaining competitive algorithms for fully dynamic 2-recoloring, and Δ -recoloring in the weighted setting remains open. Bridging the $\Theta(\log n)$ gap between the upper and lower bounds for fully dynamic 2-recoloring is a natural question. The case of fully dynamic Δ -recoloring is wide open, even in the uncapacitated setting. Another direction is the case of non-uniform capacities; we believe that Theorem 3 extends to this case.

The overprovisioned model for Δ -recoloring allows for stronger results than the resource augmented model in that the offline algorithm is permitted to utilize the same resources. Our current results assume over-provisioning in colors (with respect to the maximum degree) and capacity per color; it would be insightful to obtain results where either the capacity, or the number of colors is overprovisioned. Also, whether the $O(\log n)$ -competitive 2-recoloring result generalizes to the overprovisioned model would be interesting to resolve. Finally, can the range of Δ for our results in the over-provisioned model be improved?

2 Capacitated 2-Recoloring

2.1 Preliminaries

Let $[n]$ denote the set of integers $\{1, 2, \dots, n\}$. We denote the two clusters as C_1 and C_2 , each of which has capacity $W = \frac{1+\varepsilon}{2} \sum_{v \in V} w(v)$. A vertex v (resp. set $S \subseteq V$) is said to be *scheduled* on a cluster $C \in \{C_1, C_2\}$ if $v \in C$ (resp. $S \subseteq C$). At any point, let $N(C)$ denote the *residual capacity* of cluster $C \in \{C_1, C_2\}$ where $N(C) = W - \sum_{v \in C} w(v)$. We assume that

$w(v) \in [\frac{1}{2} \sum_{v \in V} w(v)]$ for all $v \in V$.

Our algorithms maintain a set of *bipartite connected components* \mathcal{P} induced by the request sequence. A connected component P is a maximal set of vertices such that for any $u_i, v_i \in P$, there is a path between u_i and v_i consisting only of vertices in P . Initially, $\mathcal{P} = \{\{v\} \mid v \in V\}$ so that each component consists of a single vertex. Each component $P_i \in \mathcal{P}$ for $i \in [|\mathcal{P}|]$ with bi-partition (A_i, B_i) is expressed as $P_i = (A_i, B_i)$. We refer to the weight (resp. size) of the component as $w(P_i)$ (resp. $|P_i|$), where $w(P_i) = \sum_{v \in A_i \cup B_i} w(v)$ (resp. $|P_i| = |A_i| + |B_i|$).

Component Merges. A request (u_t, v_t) between vertices in distinct components $P_1 = (A_1, B_1)$ and $P_2 = (A_2, B_2)$ leads to a merge of P_1 and P_2 into a new component P_3 , following which \mathcal{P} is updated. More precisely, if $u_t \in A_1$ (resp. B_1) and $v_t \in A_2$, then $P_3 = (A_1 \cup B_2, A_2 \cup B_1)$ (resp. $P_3 = (A_1 \cup A_2, B_1 \cup B_2)$), and analogously for the case when $v_t \in A_1$. For any component $P_i \in \mathcal{P}$, our algorithms schedule A_i and B_i in different clusters. Observe that $|\mathcal{P}| \leq n$ and $\max_{P_i \in \mathcal{P}} (|P_i|) \leq n - 1$.

Recolor Subroutine. Our algorithms employ a recoloring subroutine **Recolor**. Given a current coloring $c(v) = j$ for any vertex $v \in V$, **Recolor** (v, i) recolors v to i if the current coloring $j \neq i$ and sets $N(C_j)$ (resp. $N(C_i)$) to $N(C_j) - w(v)$ (resp. $N(C_i) + w(v)$).

Rebalancing Subroutine. Our algorithms periodically call a re-balancing subroutine **Rebalance** which takes as input a weight parameter W' , component set \mathcal{P} , and a parameter ε . It computes an *assignment* of bipartite components in \mathcal{P} to clusters such that the total weight of vertices assigned on C_1 is at most $W'(1 + \varepsilon)$. If this is not possible, it terminates. While an assignment satisfying weight constraint exactly W' can be determined

in pseudo-polynomial time $O(|\mathcal{P}|W') = O(nW')$, **Rebalance** is a FPTAS (fully polynomial-time approximation scheme) and takes $O(\frac{n^2 \ln W'}{\varepsilon})$ time. The algorithm is deferred to the full version of the paper.

2.2 An $O(n \log n)$ -Competitive Algorithm for Fully Dynamic Recoloring

We present an algorithm **Greedy-Recoloring** for fully dynamic recoloring in the unweighted setting with $(1 + \varepsilon)$ -augmentation, where $\frac{\delta}{n} \leq \varepsilon < 1$. We begin by defining a phase.

► **Definition 8.** A phase \mathcal{R} of request sequence σ is a maximal contiguous sub-sequence of σ such that, i) all components induced by requests in \mathcal{R} are bipartite and, ii) there exists a feasible assignment of components such that the total number of vertices assigned to cluster C_1 is at most $(1 + \frac{\varepsilon}{2})\frac{n}{2}$.

The sequence σ is partitioned into consecutive phases $\mathcal{R}_1, \mathcal{R}_2, \dots$ and our algorithm initializes each phase \mathcal{R}_i with a set of singleton components $\mathcal{P}_i = \{\{v\} | v \in V\}$ and assigns exactly $\frac{n}{2}$ vertices to each cluster. The set of components \mathcal{P}_i induced by requests throughout any phase \mathcal{R}_i is maintained. A phase terminates when either i) a request (u_t, v_t) is encountered where $u_t, v_t \in A_j$ or $u_t, v_t \in B_j$ for any component $P_j = (A_j, B_j)$ in \mathcal{P}_i , i.e. P_j ceases to be bipartite or ii) invoking **Rebalance** $(\mathcal{P}_i, \frac{n}{2}, \frac{\varepsilon}{2})$ returns infeasible, i.e. an assignment of components is not possible while fulfilling the capacity constraint of $(1 + \frac{\varepsilon}{2})\frac{n}{2}$. Thereafter, the next phase \mathcal{R}_{i+1} begins. The following lemma follows from the definition of a phase.

► **Lemma 9.** For an unweighted instance such that $w(v) = 1$ for all $v \in V$, OPT incurs a cost of at least 1 per phase.

Proof. If **Rebalance** does not yield an assignment of components induced by requests in phase \mathcal{R}_i , OPT must incur a cost of at least 1 during phase \mathcal{R}_i . If \mathcal{R}_i terminates because a component ceases to be bipartite, there exist requests in \mathcal{R}_i on an odd cycle for which a single recoloring is necessary. ◀

Algorithm Greedy-Recoloring. We give a high level description of our Algorithm **Greedy-Recoloring**. The pseudo code is given in the full version of the paper. Consider a request (u_t, v_t) in phase \mathcal{R}_i . If $u_t, v_t \in P$ for some component $P \in \mathcal{P}_i$ such that P stays bipartite, nothing is done. Else if P ceases to be bipartite, phase \mathcal{R}_i terminates.

If u_t, v_t are in distinct components $P_1 = (A_1, B_1), P_2 = (A_2, B_2)$ (w.l.o.g., assume that $|P_1| \geq |P_2|$), vertices in $A_1 \cup A_2$ are assigned to C_1 and vertices in $B_1 \cup B_2$ are assigned to C_2 , P_2 is merged into P_1 . If $|P_2| \leq \frac{\varepsilon n}{8}$ and recoloring is possible while fulfilling capacity constraints, P_2 is re-colored. More concretely, if $u_t \in A_1$ (resp. B_1) and $v_t \in A_2$ (resp. B_2), vertices in B_2 (resp. A_2) are recolored and assigned to C_1 (resp. C_2), and vertices in A_2 (resp. B_1) are recolored and assigned to C_2 (resp. C_1) if residual capacities $N(C_1), N(C_2)$ are sufficient. Otherwise, **Rebalance** $(\mathcal{P}_i, \frac{n}{2}, \frac{\varepsilon}{2})$ is invoked. If **Rebalance** yields a feasible component assignment, vertices are appropriately recolored such that $N(C_1)$ and $N(C_2)$ are both at least $\frac{\varepsilon}{2} \cdot \frac{n}{2} = \frac{\varepsilon n}{4}$.

Proof of Theorem 1. By virtue of the recolorings performed, Algorithm **Greedy-Recoloring** always maintains a 2-coloring of the graph induced by the sequence of requests in a phase. Note that **Rebalance** incurs a total recoloring cost at most n . Before the first request arrives, we call **Greedy-Recoloring** $(\{v | v \in V\}, \frac{n}{2})$. For the sake of analysis, we say a component $P_j \in \mathcal{P}_i$ during phase \mathcal{R}_i is *small* if $|P_j| \leq \frac{\varepsilon n}{8}$ and *large* otherwise. Note that the algorithm always recolors the smaller component, P_2 when $P_1 = (A_1, B_1)$ and $P_2 = (A_2, B_2)$ are

merged, as long as component P_2 is small and capacity constraints are *not violated*. In this case, $|P_1| \geq 2|P_2|$, and each vertex in P_2 is charged a recoloring cost of 1. If P_2 is large, **Rebalance** is called.

The key observation is that whenever a small component P_1 is created as a result of two small components P_1 and P_2 being merged, and A_2 and B_2 cannot be recolored because either $N(C_1) - |A_2| < |B_2|$ or $N(C_2) - |B_2| < |A_2|$, at least $\frac{\varepsilon n}{8}$ vertices (let S denote the set of these vertices) must have been successfully recolored since the last time **Rebalance** was called. This holds since $(1 + \varepsilon)\frac{n}{2} - (1 + \frac{\varepsilon}{2})\frac{n}{2} - |P_1| \geq \frac{\varepsilon n}{8}$. We charge the recoloring cost of n incurred as a result of **Rebalance** uniformly to all vertices in S , i.e. every vertex in S receives a charge of $\frac{n}{\varepsilon n/8} = \frac{8}{\varepsilon}$. Each time a vertex is charged in this manner, the component containing it doubles in size. Hence, any vertex is charged at most $\frac{8}{\varepsilon} \log \frac{\varepsilon n}{8} = O(\log n)$ while it is part of a small component.

On the other hand if P_2 is large, **Rebalance** is called. The recoloring cost is charged uniformly to all vertices in P_2 , so that each vertex in P_2 receives a charge of at least $\frac{n}{\varepsilon n/8} = \frac{8}{\varepsilon}$. Thus, every time a vertex is part of a large component it can be charged a recoloring cost at most $\log(\frac{8}{\varepsilon}) = O(1)$ since its component size doubles and the maximum size of any component is n .

If recoloring does not lead to a call to **Rebalance**, vertices in P_2 are simply charged 1 for the recoloring. Thus, every vertex receives a total charge of $\log n + \frac{8}{\varepsilon} \log(\frac{8}{\varepsilon}) + \frac{8}{\varepsilon} \log \frac{\varepsilon n}{4} = O(\log n)$ throughout the phase. The total recoloring cost incurred by the algorithm is $O(n \log n)$, and by Lemma 9, this yields $O(n \log n)$ -competitiveness. ◀

2.3 An $O(\log n)$ -Competitive Algorithm for Online Recoloring

In this section, we present our main result for 2-recoloring which is an asymptotically optimal $O(\log n)$ -competitive algorithm. Recall that in the uniform capacity setting, the graph induced by request sequence σ admits a proper coloring such that the weight of vertices assigned any color is exactly $\frac{1}{2} \sum_{v \in V} w(v)$.

Notation and Definitions. Given an initial coloring c_0 , an arbitrary coloring c and $S \subseteq V$, let $d_S(c, c_0)$ denote the weighted Hamming distance of c from c_0 restricted to S , i.e. $d_S(c, c_0) = \sum_{v \in S: c_0(v) \neq c(v)} w(v)$. Let $d_S(c) := d_S(c, c_0)$ for brevity. For a bipartite component $P_i = (A_i, B_i)$, there are exactly two feasible colorings; let c_1 denote the coloring for which vertices in A_i (resp. B_i) are colored 1 (resp. 2) and c_2 denote the coloring for which vertices in A_i (resp. B_i) are colored 2 (resp. 1). Let c_m denote the optimal coloring where $m \in \{1, 2\}$, such that $d_{P_i}(c_m) = \min\{d_{P_i}(c_1), d_{P_i}(c_2)\}$, i.e. c_m corresponds to a coloring incurring the minimum recoloring cost for P_i w.r.t. the initial coloring c_0 . Since there are only two possible colorings, it follows that $d_{P_i}(c_m) \leq \frac{w(P_i)}{2}$. Our algorithm assigns vertices in A_i on cluster C_m and vertices in B_i on cluster C_{3-m} . Finally, let $E(P_i)$ denote the *estimated weight* of component P_i . Our algorithm maintains the property that $E(P_i) \geq \frac{w(P_i)}{(1 + \frac{\varepsilon}{4})}$, and periodically recomputes optimal colorings whenever the total weight $w(P_i) > (1 + \frac{\varepsilon}{4})E(P_i)$, and updates $E(P_i)$ to $w(P_i)$. Let $c(S)$ denote the restriction of some coloring c to set $S \subseteq V$.

Algorithm Follow-Greedy. The algorithm begins with the set of components $\mathcal{P} = \{\{v\} | v \in V\}$ and $E(\{v\}) = w(v)$ for all v . For any request (u_t, v_t) such that $u_t \in P_1, v_t \in P_2$ where w.l.o.g., $w(P_1) \geq w(P_2)$, the algorithm merges P_2 into P_1 . There are two cases to consider: i) if $c(u_t) = c(v_t)$ and $w(P_1) \leq E(P_1)(1 + \frac{\varepsilon}{4})$ following the merge and, ii) $w(P_1) > E(P_1)(1 + \frac{\varepsilon}{4})$ following the merge. In the former case, only vertices in P_2 are recolored as long as capacity

95:10 Competitive Capacitated Online Recoloring

constraints are respected. In the latter case, the algorithm computes an optimal coloring c_m for P_1 . If $c_m(P_1) \neq c(P_1)$, where c is the current coloring and recoloring P_1 according to c_m respects capacity constraints, P_1 is recolored. If capacity constraints are violated at any point, **Follow-Greedy** calls (and transitions to) algorithm **Greedy-Recoloring**(\mathcal{P}, W). The complete pseudo-code of the algorithm is deferred to the full version of the paper.

Analysis. Whenever components are *greedily* assigned an optimal coloring such that the total weight of components exceeds W , **Follow-Greedy** transitions to **Greedy-Recoloring**. At this point, we show that $OPT_\sigma = \Omega(W)$, and bound the recoloring cost incurred by **Greedy-Recoloring** on the request sequence thereafter by $O(W \log n)$. On the other hand, we show that the total recoloring cost incurred for any component P whenever $w(P)$ increases (before **Greedy-Recoloring** is ever called) can be bounded by $O(\log n)$ times the recoloring cost incurred by **OPT** on P . Combining the two bounds yields $O(\log n)$ -competitiveness.

Lemma 10 establishes a monotonicity property— distance to optimal colorings for any component P is monotonically non-decreasing with its weight $w(P)$.

► **Lemma 10.** *Let P and P' be components at times t_1 and t_2 respectively, where $1 \leq t_1 \leq t_2 \leq T$ and $P \subseteq P'$. If c_m and c'_m are optimal colorings for P and P' respectively, then $d_P(c_m) \leq d_{P'}(c'_m) \leq d_{P'}(c'_m)$.*

Proof. We note that if $c_m(P) \neq c'_m(P)$, then

$$d_P(c_m) = \sum_{v \in P: c_m(v) \neq c_0(v)} w(v) \leq \sum_{v \in P: c'_m(v) \neq c_0(v)} w(v) = d_P(c'_m)$$

since c_m is an optimal coloring for vertices in P . Using this we have,

$$\begin{aligned} d_{P'}(c'_m) &= \sum_{v \in P: c'_m(v) \neq c_0(v)} w(v) + \sum_{v \in P' \setminus P: c'_m(v) \neq c_0(v)} w(v) \geq \sum_{v \in P: c'_m(v) \neq c_0(v)} w(v) \\ &= d_P(c'_m). \end{aligned} \quad \blacktriangleleft$$

Our algorithm periodically recomputes an optimal coloring c_m for any component when its weight increases by least a $(1 + \frac{\varepsilon}{4})$ factor, and re-colors vertices according to c_m (if needed). Let T denote the time step before **Greedy-Recoloring** is called; if **Greedy-Recoloring** is never called, $T = |\sigma|$. A component $P_i \in \mathcal{P}$ is said to be *alive* during a time interval $T_{P_i} = [1, T']$ where $T' \leq T$, if for all requests $\sigma_t = (u_t, v_t) \in \sigma_{\leq T'}$ such that $u_t \in P_i, v_t \in P_j$, $w(P_i) \geq w(P_j)$; i.e. P_i is not the *smaller* of the two components merged during time $t \in T_{P_i}$.

► **Lemma 11.** *Let P_i be an alive component during interval $T_{P_i} = [1, T']$ and c_m^t denote an optimal coloring for P_i at time t . For all $t \in T_{P_i}$ the coloring c maintained by **Follow-Greedy** satisfies $d_{P_i}(c) \leq d_{P_i}(c_m^t) + \frac{\varepsilon}{4}w(P_i)$, where $w(P_i)$ is the weight of P_i at time t .*

Proof. Let $c(P_i)$ be the coloring maintained by the algorithm. Let $t_a, t_b \in T_{P_i}$ be time-steps such that $t_a < t_b$ at which the algorithm re-computes an optimal coloring for P_i and recolors P_i if necessary. Thus, $c(P_i) = c_m^t(P_i)$ at $t = t_a$ and $t = t_b$ respectively. For any $t \in (t_a, t_b)$, if $c = c_m^t$, then $d_{P_i}(c) = d_{P_i}(c_m^t)$. For all other t , $d_{P_i}(c) \geq d_{P_i}(c_m^t)$ by Lemma 10. Let S denote the set of vertices which are merged to P_i in the time interval $(t_a + 1, t_b)$. Since $w(S) \leq \frac{\varepsilon}{4}w(P_i)$, $d_{P_i}(c) \leq d_{P_i}(c_m^t) + \frac{\varepsilon}{4}w(P_i) \leq d_{P_i}(c_m^t) + \frac{\varepsilon}{4}w(P_i)$, completing the proof. ◀

Lemma 12 gives a lower bound on OPT_σ if **Greedy-Recoloring** is called.

► **Lemma 12.** *If **Follow-Greedy** calls **Greedy-Recoloring** at time t , then $OPT_{\sigma_{\leq t}} = \Omega(W)$.*

Proof. Note that the residual capacity $N(C_i) = (1 + \varepsilon)W - \sum_{v \in C_i} w(v)$ for any $i \in \{1, 2\}$. Algorithm **Greedy-Recoloring** is called whenever following a merge of P_2 into P_1 , either i) an optimal coloring for P_1 violates capacity constraints or ii) recoloring vertices in P_2 violates capacity constraints. Let c denote the *infeasible* coloring in any case, s.t. i) $c = c_m$ (the optimal coloring) when $w(P_1) > E(P_1)(1 + \frac{\varepsilon}{4})$ or, ii) $c(v) = 1$ for all $v \in B_2$, and $c(v) = 2$ for all $v \in A_2$ (see lines 13-15, 20 of **Follow-Greedy**), such that $\sum_{P_i \in \mathcal{P}} \sum_{v \in P_i: c(v)=j} w(v) > W(1 + \varepsilon)$ for some $j \in \{1, 2\}$. Note that such a j must exist since the capacity constraint is violated for one of the two clusters by the coloring c . By Lemma 11, it follows that $d_{P_i}(c) \leq d_{P_i}(c_m) + \frac{\varepsilon}{4}w(P_i)$ where c_m denotes an optimal coloring at time t for any component P_i . Note that $OPT_\sigma \geq \sum_{P_i \in \mathcal{P}} d_{P_i}(c_m)$ since this lower bound disregards capacity constraints. Thus, $\sum_{P_i \in \mathcal{P}} d_{P_i}(c) \leq \sum_{P_i \in \mathcal{P}} [d_{P_i}(c_m) + \frac{\varepsilon}{4}w(P_i)] \leq OPT_\sigma + \frac{\varepsilon}{2}W$ since $\sum_{P_i \in \mathcal{P}} w(P_i) = 2W$. Thus, $OPT_\sigma \geq \sum_{P_i \in \mathcal{P}} d_{P_i}(c) - \frac{\varepsilon}{2}W$. Now, note that $\sum_{P_i \in \mathcal{P}} d_{P_i}(c) \geq \varepsilon W$ since cluster capacities have been exceeded by at least εW , which is only possible if at least a set S of vertices with weight at least εW change their initial coloring. Thus, $OPT_\sigma \geq \varepsilon W - \frac{\varepsilon}{2}W = \frac{\varepsilon}{2}W = \Omega(W)$. \blacktriangleleft

Lemma 13 bounds the recoloring cost of **Greedy-Recoloring**. The proof is similar to the proof of Theorem 1 and deferred to the full version of the paper.

► **Lemma 13.** *The total recoloring cost incurred by Greedy-Recoloring when it is called by Follow-Greedy is $O(W \log n)$.*

► **Lemma 14.** *Follow-Greedy incurs a total recoloring cost of $O(OPT_\sigma \log n)$ before Greedy-Recoloring is ever called.*

Proof. Let $T = |\sigma|$ if **Greedy-Recoloring** is never called; otherwise $T = t - 1$ if it is called at time t . We analyze the cost of **Follow-Greedy** on the set of all *alive* components \mathcal{P}_A until time T . Any alive component P' at any point satisfies $P' \subseteq P$ for some alive component $P \in \mathcal{P}_A$. Let $OPT(P)$ denote the cost incurred by OPT on P until time T . We show that the total cost of **Follow-Greedy** on P is bounded by $OPT(P)O(\log n)$ to conclude the proof.

For any component $P \in \mathcal{P}_A$, let $I_P = \{[0 = t_0, t_1], [t_1, t_2], [t_2, t_3], \dots, [t_f, t_{f+1} = T]\}$ denote the partition of time interval $[0, T]$ such that, for all $j \in [f]$ i) $P_j \subseteq P$ is a sub-component of P for which an optimal recoloring is computed at time t_j , ii) $(1 + \frac{\varepsilon}{4})w(P_{j-1}) \leq w(P_j)$ and iii) $P_1 \subseteq P_2 \subseteq \dots \subseteq P_{f+1} = P$. Recall that optimal colorings are only computed whenever component weights increase by at least a $(1 + \frac{\varepsilon}{4})$ factor. Let \mathcal{Q}_j denote the set of components that are merged to P_{j-1} during time interval $[t_{j-1}, t_j]$ to yield P_j . Note that for components $Q \in \mathcal{Q}_j$, $w(Q) \leq w(P_{j-1}) \leq (1 + \frac{\varepsilon}{4})w(P_{j-1})$ by definition of P_j and \mathcal{Q}_j and the aforementioned property ii). Thus, $w(P_j) \leq 2(1 + \frac{\varepsilon}{4})w(P_{j-1})$.

There are two types of recolorings performed during time interval $[t_{j-1}, t_j]$:

- (*Type I*) P_j is recolored at time t_j .
- (*Type II*) $Q \in \mathcal{Q}_j$ is recolored at $t_Q \in [t_{j-1}, t_j]$ if $c(u_t) = c(v_t)$ where $u_t \in P_{j-1}$ and $v_t \in Q$.

We charge the cost of both Type I and Type II colorings to OPT . For Type I colorings, note that $c_m^{t_{j-1}}(P_{j-1}) \neq c_m^{t_j}(P_j)$, and hence $OPT(P) \geq OPT(P_j) \geq \frac{1}{2}w(P_j) \geq \frac{1}{2} \frac{w(P_{j-1})}{2(1 + \frac{\varepsilon}{4})} = \frac{w(P_j)}{4 + \varepsilon}$ by Lemma 10 and the upper bound on $w(P_j)$. We charge every $v \in P_j$ s.t. $w(v) > \frac{w(P_j)}{2n}$ an amount $2w(v)$ to reflect this cost. Let $j_{max} \leq f$ denote the largest j for which P_j is recolored in this manner. Since these recolorings may happen every time a component undergoes a $(1 + \frac{\varepsilon}{4})$ factor increase in weight, each vertex in $P_{j_{max}}$ is charged at most $O(\log_{(1 + \frac{\varepsilon}{4})} n) = O(\log n)$ Type I recoloring costs. Accounting for all $P \in \mathcal{P}_A$, the total Type I recoloring cost is $O(OPT_\sigma \log n)$.

For Type II recolorings, we consider two cases depending on whether P_j is recolored at time t_j or not. Consider a component $Q \in \mathcal{Q}_j$ that is merged to P_{j-1} at time t_Q for $j \in [f+1]$ and recolored. If P_j is recolored at time t_j , the Type II recoloring cost for Q can be absorbed in the Type I recoloring cost for P_j by charging each vertex $v \in P_j$ where $w(v) > \frac{w(P_j)}{2n}$ a cost $4w(v)$ (instead of $2w(v)$ as before).

On the other hand, if P_j is not recolored at time t_j (and hence the cost of a Type II recoloring cannot be associated to a Type I recoloring), we charge the cost as follows. Let $\mathcal{Q}'_j \subseteq \mathcal{Q}_j$ denote the set of components that are recolored by **Follow-Greedy** and consider a component $Q \in \mathcal{Q}'_j$ that is recolored at time t_Q . By Lemma 11, it follows that for all $t \in [t_{j-1}, t_Q)$, the coloring c maintained by **Follow-Greedy** satisfies $d_Q(c) \leq d_Q(c_m^t) + \frac{\varepsilon}{4}w(Q)$, where c_m^t corresponds to an optimal coloring for Q at time t . Moreover, $OPT(P_j) \geq d_{P_{j-1}}(c_m^{t_j-1}) + \sum_{Q \in \mathcal{Q}'_j} \frac{1}{2} \frac{w(Q)}{(1+\frac{\varepsilon}{4})} = d_{P_{j-1}}(c_m^{t_j-1}) + \sum_{Q \in \mathcal{Q}'_j} \frac{w(Q)}{(2+\frac{\varepsilon}{2})}$ since P_j is not recolored. In other words, the optimal coloring for P_j is sub-optimal on a set $S \subseteq Q$ of vertices of weight at least $\frac{w(Q)}{(1+\frac{\varepsilon}{4})}$ and hence, incurs a recoloring cost of at least $\frac{1}{2}w(S)$ for recoloring Q . We charge every vertex $v \in Q$ for which $w(v) > \frac{w(Q)}{2n}$ a cost $2w(v)$. For any j and $Q \in \mathcal{Q}'_j$ for which Type II recolorings happen, every vertex v in Q can only be charged $O(\log n)$ times since the weight of the component containing v doubles each time v is charged. This yields a total cost of $O(OPT_\sigma \log n)$ for Type II recolorings, completing the proof. \blacktriangleleft

Proof of Theorem 3. We prove that **Follow-Greedy** is $O(\log n)$ -competitive. If **Greedy-Recoloring** is not invoked, Lemma 14 yields $O(\log n)$ competitiveness. Else, by Lemmas 14 and 13 a total cost of $O(OPT_\sigma \log n + W \log n) = O(W \log n)$ is incurred. Combined with Lemma 12, this yields $O(\log n)$ -competitiveness. \blacktriangleleft

3 Capacitated Online Δ -recoloring

We give deterministic and randomized algorithms for capacitated Δ -recoloring in an overprovisioned setting. Formally, we have n vertices, Δ colors, each with capacity $(1+\varepsilon)\frac{n}{\Delta}$, such that the maximum degree of the graph induced by request sequence σ is at most $(1-\varepsilon)\Delta$. We analyze competitiveness with respect to an optimal offline algorithm with *no capacity constraints*. Initially all vertices are assigned to Δ colors such that each color has exactly $\frac{n}{\Delta}$ vertices (for convenience, we assume that n is divisible by Δ).

Our algorithms maintain a list of feasible colors $L(v) \subseteq [1, \Delta]$ for each vertex $v \in V$. We derive a lower bound for OPT_σ (similarly to [6] for uncapacitated $(\Delta+1)$ -recoloring) by considering the size of the minimum vertex cover of graph G_M induced by the set of monochromatic edges (u_t, v_t) (with respect to the initial coloring c_0), i.e. $c_0(u_t) = c_0(v_t)$. The following lemma is immediate since OPT must recolor at least one endpoint of a monochromatic edge in G_M .

\blacktriangleright **Lemma 15** ([6]). $OPT_\sigma \geq |C^*|$, where C^* denotes a minimum vertex cover of G_M .

Lemma 15 holds for the uncapacitated case [6]; thus, our algorithms in the overprovisioned setting are competitive with respect to an optimal algorithm which is not capacity constrained.

Maintaining a 2-approximate vertex cover. Both of our algorithms maintain a 2-approximate vertex cover C of G_M *online*, i.e., $|C| \leq 2|C^*|$ using a simple greedy algorithm: on arrival of an edge (u_t, v_t) such that $c_0(u_t) = c_0(v)$, if neither of u_t or v_t are in C , both u_t and v_t are added to C . Consider any distinct pairs of vertices $(u_i, v_i), (u_j, v_j)$ added to C . By virtue of the algorithm, vertices u_i and v_i can not be adjacent to u_j or v_j in G_M . Thus, any vertex cover must include one of u_t or v_t whenever u_t, v_t are added by the greedy algorithm after a request (u_t, v_t) is encountered. It follows that $C \leq 2|C^*|$.

The Algorithm. We give a generic algorithm Δ -Recoloring below that invokes a Recolor subroutine. Our deterministic and randomized algorithms differ in their respective implementations of this subroutine. The vertex cover C is initialized to \emptyset . On any request (u_t, v_t) , if $c(u_t) = c(v_t)$ there are multiple cases to consider:

1. If $u_t, v_t \notin C$: add u_t, v_t to C , and call $\text{Recolor}(u_t)$.
2. Elsf $u_t \in C, v_t \notin C$: call $\text{Recolor}(u_t)$.
3. Else: if both $u_t, v_t \in C$, call $\text{Recolor}(\arg \max_{w \in \{u_t, v_t\}} \deg(w))$.

3.1 A Deterministic $O(\Delta)$ -Competitive Algorithm

We give a deterministic algorithm by describing a *deterministic* Recolor subroutine which is used in conjunction with algorithm Δ -Recoloring. As before, let $N(C_i) = (1 + \varepsilon)\frac{n}{\Delta} - \sum_{v \in C_i} w(v) = (1 + \varepsilon)\frac{n}{\Delta} - |C_i|$ denote the residual capacity of cluster C_i , for any i .

Deterministic Recolor(v). Suppose v is currently assigned to color i . Among the list of feasible colors $L(v)$ for v , we recolor v with color j with maximum residual capacity; i.e., $C_j = \arg \max_{\{C_i: i \in L(v), N(C_i) > 0\}} N(C_i)$. The variables $N(C_i)$ (resp. $N(C_j)$) are incremented (resp. decremented) by 1, and the list $L(v)$ updated. If $N(C_j) = 0$ for all $j \in L(v)$, then the following Rebalance procedure is called and all colors are assigned at most $\lceil \frac{n}{\Delta} \rceil$ vertices.

Deterministic Rebalance. For any graph G with maximum degree at most r , an equitable coloring is a proper coloring such that for any $i, j \in [r + 1]$ the number of vertices assigned color i differs by the number of vertices assigned color j by at most 1. Kierstead et al. [19] present an $O(rn^2)$ -time algorithm to compute a $(r + 1)$ -equitable coloring. We use their algorithm as the rebalancing procedure to recompute a coloring for the subgraph G_t induced by $\sigma_{\leq t}$. The cost of Rebalance is trivially bounded by $O(n)$.

Proof of Theorem 6. Let $|C^*|$ denote the *size* of the optimal vertex cover for the graph G_M . Our algorithm only recolors vertices in C . Since the maximum degree is $(1 - \varepsilon)\Delta$, there are at least $\varepsilon\Delta$ feasible colors for each vertex when recoloring happens. If there exists a feasible color $j \in L(v)$ s.t. $N(C_j) \geq 1$, a cost of 1 is incurred. Thus, if Rebalance is never called throughout the request sequence, the total cost is at most $O(2 \cdot |C^*| \cdot (1 - \varepsilon)\Delta) = O(|C^*|\Delta)$. Combined with Lemma 15, this yields $O(\Delta)$ -competitiveness.

Now, let v denote the first such vertex for which a rebalancing procedure is called. This implies that for at least $\varepsilon\Delta$ feasible colors $j \in L(v)$, s.t. $N(C_j) = 0$. It follows that $\text{OPT} \geq |C^*| \geq \varepsilon\Delta(\frac{\varepsilon n}{\Delta} - 1) = \varepsilon^2 n - \varepsilon\Delta \geq \frac{\varepsilon^2 n}{2}$ in this case, since only vertices in C are recolored. Since the sum of degrees of all nodes in the subgraph G_T induced by σ is $n\Delta(1 - \varepsilon)$, the re-balancing procedure can be called $O(\frac{n\Delta(1-\varepsilon)}{\varepsilon^2 n}) = O(\Delta)$ times, incurring a total re-balancing cost $O(n\Delta)$ since a single call to re-balancing incurs $O(n)$ recolorings. The total number of recolorings between re-balancing calls is bounded by $O(n\Delta)$. Since $\text{OPT} = \Omega(n)$ in this case, the algorithm is $O(\Delta)$ -competitive, completing the proof. \blacktriangleleft

3.2 A Randomized $O(1)$ Competitive Algorithm

Our randomized algorithm uses *randomized* Recolor and Rebalance subroutines in conjunction with algorithm Δ -Recoloring.

Randomized Recolor(v). Suppose v is currently assigned color i . We pick a feasible color j for v uniformly at random from $L(v)$. If $N(C_j) \geq 1$, recolor v to j and assign to C_j , and increment $N_C(i)$ (resp. decrement $N_C(j)$) by 1; otherwise, call randomized Rebalance.

Randomized Rebalance. Initially, for each vertex v , the feasible list of colors is $L(v) = [\Delta]$. Thereafter, a coloring for V is sequentially determined as follows: for each $v \in V$, a random color c is picked from the current list of feasible colors $L(v)$, and assigned to v . For all neighbors u of v , c is removed from $L(u)$ and the process repeats.

The following holds since the maximum degree of any vertex is always bounded by $(1 - \varepsilon)\Delta$.

► **Observation 16.** $|L(v)| \geq \varepsilon\Delta$ for any $v \in V$ and any time t .

Analysis of Rebalance. A key property of randomized rebalancing is that the distribution of the vertices across the colors is very close to balanced. This is formalized by the following lemma, whose proof utilizes a martingale argument and Azuma's inequality.

► **Lemma 17.** *Randomized Rebalance computes a proper coloring c satisfying i) the number of vertices assigned any color i where $i \in [\Delta]$ is at most $(1 + \frac{\varepsilon}{2})\frac{n}{\Delta}$ with probability at least $1 - \Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$ and ii) the probability that any two vertices u, v have the same color is at most $\frac{1}{\varepsilon\Delta}$.*

Proof. We first note that rebalancing always computes a proper coloring. To establish a high probability upper bound on the number of vertices in any color, we use a martingale argument. Let X_1, X_2, \dots, X_n denote i.i.d. random variables uniformly generated from the interval $[0, 1]$ in advance. In the randomized **Rebalance** subroutine, we can view the color assignment for the j th vertex v as follows: the algorithm selects the k th color in $L(v)$ for $k = \lceil X_j \cdot |L(v)| \rceil$. This ensures that any color $c \in L(v)$ is picked with probability $\frac{1}{|L(v)|}$. This view allows us to condition our martingale on independent random variables X_1, \dots, X_i instead of the actual color choices, which are not independent in general.

Let Y_i denote the number of vertices assigned color i once rebalancing is completed. Note that Y_i depends on random variables X_1, \dots, X_n all of which are mutually independent. Let Z denote the Doob martingale where $Z_j = \mathbb{E}[Y_i | X_1, \dots, X_j]$, i.e. Z_j is the expected value of Y_i given the random choices of the first j vertices, and let $Z_0 = \mathbb{E}[Y]$.

To compute a high probability bound on the value of Z_n , we first compute $\mathbb{E}[Z_0] = \mathbb{E}[Y_i]$. We note that the coloring computed by randomized **Rebalance** is symmetric with respect to any color. Thus, $\mathbb{E}[Y_i] = \mathbb{E}[Y_j]$ for any $i \neq j$ where $i, j \in [\Delta]$. Since $\mathbb{E}[\sum_{i \in [\Delta]} Y_i] = n$, we have that $\mathbb{E}[Y_i] = \frac{n}{\Delta}$ for any $i \in \Delta$. Moreover, note that $Z_j - Z_{j-1} \leq 1$ for any $j \in [n]$. We apply the Azuma-Hoeffding inequality (e.g., see Theorem 13.6 in [20]) to derive

$$\Pr \left[Z_n - Z_0 \geq \frac{\varepsilon n}{2\Delta} \right] = \Pr \left[Z_n - \frac{n}{\Delta} \geq \frac{\varepsilon n}{2\Delta} \right] \leq e^{-\frac{2\varepsilon^2 n^2}{4\Delta^2 n}} = e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$$

Taking a union bound for all Δ colors, we have that the probability that for any color i , the number of vertices assigned color i is at most $(1 + \frac{\varepsilon}{2})\frac{n}{\Delta}$ is at least $1 - \Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$. This completes the proof for property i) of coloring c given by randomized rebalancing.

For property ii), let v be colored after u when randomized rebalancing is called. Note that v has at least $\frac{1}{\varepsilon\Delta}$ feasible colors in $L(v)$ at the time it is colored (by Observation 16) and chooses a color randomly. Since u is a non-neighbor, u can be assigned one of the colors in $L(v)$; since v is assigned a random color in $L(v)$, we obtain $\Pr[c(u) = c(v)] \leq \frac{1}{\varepsilon\Delta}$. ◀

Analysis of recoloring. Our algorithm recolors only when $c(u_t) = c(v_t)$ for a request (u_t, v_t) at any time t . We first place an upper bound on the probability of recoloring. The proof of the following lemma is deferred to the full version of the paper.

► **Lemma 18.** *For a request (u_t, v_t) in which at least one of u_t or v_t has been recolored before time t , $\Pr[c(u_t) = c(v_t)] \leq \frac{1}{\varepsilon\Delta}$.*

For our analysis, we partition the request sequence into phases where a phase ends whenever a rebalancing occurs. This yields a (possibly empty) sequence of complete phases followed by an incomplete phase. Note that the length of any complete phase is a random variable. We first show that each complete phase has $\Omega(n)$ recolorings with high probability.

► **Lemma 19.** *A complete phase has at least $\varepsilon^2 n/4$ recolorings with probability at least $1 - 2\Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$.*

Proof. We establish the claim by showing that if at most $\frac{\varepsilon^2 n}{4}$ recolorings take place in phase, then the number of vertices recolored to any particular color j is at most $(1 + \varepsilon)\frac{n}{\Delta}$ with high probability, indicating that the phase has not completed. In the remainder of the proof, we assume that $t = \frac{\varepsilon^2 n}{4}$ recolorings take place. Fix a color j . Let Y_i denote the random variable which is 1 if there exists a vertex v such that v is recolored to j during its i^{th} recoloring. Note that $\Pr[Y_i = 1 | Y_{i-1}, \dots, Y_1] \leq \frac{1}{\varepsilon\Delta}$ by Observation 16. We are interested in the random variable $Y = \sum_{i=1}^t Y_i$, and note that the total number of vertices assigned to color j is bounded by $(1 + \frac{\varepsilon}{2})\frac{n}{\Delta}$ at the beginning of a phase with high probability, by Lemma 17. Moreover, Y is stochastically dominated by $Z = \sum_{i=1}^t Z_i$, where Z_1, Z_2, \dots, Z_t are i.i.d. Bernoulli random variables such that $Z_i = 1$ with probability $\frac{1}{\varepsilon\Delta}$ and 0 otherwise. Note that $\mathbb{E}[Z] \leq \frac{\varepsilon n}{4\Delta}$. Thus, to upper bound Y , we can use a Chernoff bound (see p. 69 of [20]) for Z :

$$\Pr[Y \geq \frac{\varepsilon n}{2\Delta}] \leq \Pr[Z \geq \frac{\varepsilon n}{2\Delta}] \leq \Pr[Z \geq 2\mathbb{E}[Z]] \leq e^{-\frac{\varepsilon n}{12\Delta}}$$

Taking a union bound over all j , we have that the probability the maximum number of vertices assigned any color $j \in \Delta$ is at most $(1 + \varepsilon)\frac{n}{\Delta}$ with probability at least $1 - \Delta e^{-\frac{\varepsilon n}{12\Delta}}$. Adding the preceding failure probability to the failure probability of Lemma 17 and noting that $\Delta e^{-\frac{\varepsilon n}{12\Delta}} \leq \Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$ (for $\Delta \geq 6/\varepsilon$) yields the desired claim. ◀

Lemma 20 bounds the number of new edges requested in any completed phase by $\Omega(n\Delta)$ with high probability. This bounds the number of phases by $O(1)$.

► **Lemma 20.** *The number of new requests (u_t, v_t) in any complete phase such that (u_t, v_t) is requested for the first time at time t is at least $\frac{\varepsilon^3 n\Delta}{4(1+\varepsilon)}$ with probability at least $1 - 3\Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$.*

Proof. Let Y_i denote the random variable which is 1 if the i^{th} newly requested edge (u_i, v_i) during any phase leads to a recoloring by the algorithm, and 0 otherwise. Let $N = \frac{\varepsilon^3 n\Delta}{4(1+\varepsilon)}$. By Lemma 18, $\Pr[Y_i = 1 | Y_{i-1}, \dots, Y_1] \leq \frac{1}{\varepsilon\Delta}$. We use the fact that $Y = \sum_{i=1}^N Y_i$ is stochastically dominated by $Z = \sum_{i=1}^N D_i$ where D_i is a Bernoulli random variable that is 1 with probability $\frac{1}{\varepsilon\Delta}$ and 0 otherwise. Thus, we can apply a Chernoff bound (see p. 69 of [20]) to Z to get a high probability bound on Y . We analyze the probability that $Z = \sum_{i=1}^N D_i \geq \frac{\varepsilon^2 n}{4}$, i.e. the probability that the number of recolorings is at least $\frac{\varepsilon^2 n}{4}$ recolorings for N new requests in the phase. Then $\mathbb{E}[Z] = \frac{\varepsilon^3 n\Delta}{4\varepsilon(1+\varepsilon)\Delta} = \frac{\varepsilon^2 n}{4(1+\varepsilon)}$. We derive

$$\Pr\left[Y \geq \frac{\varepsilon^2 n}{4}\right] \leq \Pr\left[Z \geq \frac{\varepsilon^2 n}{4}\right] = \Pr\left[Z \geq (1 + \varepsilon)\mathbb{E}[Z]\right] \leq e^{-\frac{\varepsilon^4 n}{12(1+\varepsilon)}}.$$

We obtain the desired claim by adding the above failure probability to that of Lemma 19 and noting that Δ is sufficiently large ($\geq 6/\varepsilon^2$). ◀

Lemma 21 bounds the expected number of recolorings throughout the algorithm.

► **Lemma 21.** *The expected number of recolorings performed during the first phase is $O(|C^*|)$. The expected number of recolorings of the algorithm is $O(n)$.*

Proof. From Lemma 18, it follows that at any time t and any vertex v such that u is not a neighbor of v at time t , $\Pr[c(u) = c(v)] \leq \frac{\varepsilon}{\Delta}$. Since only vertices in the vertex cover \mathcal{C} are recolored, the number of events for which a recoloring may happen is $|\mathcal{C}|\Delta(1 - \varepsilon)$. Each such event happens with probability at most $\frac{1}{\varepsilon\Delta}$ and since $|\mathcal{C}| \leq 2|C^*|$, it follows that the expected number of recolorings by the algorithm is bounded by $2|C^*|\frac{(1-\varepsilon)}{\varepsilon} = O(|C^*|)$.

To establish the second statement of the lemma, we use Lemma 18 and the fact that the total number of distinct edges is bounded by $n\Delta$ to derive that the expected total number of recolorings of the algorithm is at most $n + n\Delta/(\varepsilon\Delta) = n(1 + 1/\varepsilon)$. ◀

Finally, combining the lemmas above yields $O(1)$ -competitiveness for our randomized algorithm, thus proving Theorem 7.

Proof of Theorem 7. By Lemma 15, $OPT_\sigma \geq |C^*|$. We consider two cases. First, if $|C^*| \leq \frac{\varepsilon^2 n}{4(1+\varepsilon)}$, then the number of edges introduced is at most $\frac{\varepsilon^3 n \Delta}{4(1+\varepsilon)}$. By Lemma 20, with probability at least $1 - 3\Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$, the first phase does not complete. The failure probability is $1/\text{poly}(n)$ for $\Delta = O(\sqrt{n/\log n})$. Therefore, by Lemma 21, the expected cost of the algorithm is $O(|C^*|) + n\Delta/\text{poly}(n) = O(|C^*|)$.

Second, if $|C^*| > \frac{\varepsilon^2 n}{4(1+\varepsilon)}$, then by Lemma 21, the expected cost of all the recolorings is $O(n)$. By Lemma 20, the number of phases is $O(1)$ with probability at least $1 - 3\Delta e^{-\frac{\varepsilon^2 n}{2\Delta^2}}$. Therefore, the expected rebalancing cost is at most $(1 - 1/\text{poly}(n))O(n) + n\Delta/\text{poly}(n) = O(n)$. Therefore, the expected total cost is $O(|C^*|)$, thus completing the proof. ◀

References

- 1 Zaid Allybokus, Nancy Perrot, Jérémie Leguay, Lorenzo Maggi, and Eric Gourdin. Virtual function placement for service chaining with partial orders and anti-affinity rules. *Networks*, 71(2):97–106, 2018.
- 2 The Kubernetes Authors. Assigning pods to nodes, 2024. URL: <https://kubernetes.io/docs/concepts/scheduling-eviction/assign-pod-node/>.
- 3 Chen Avin, Marcin Bienkowski, Andreas Loukas, Maciej Pacut, and Stefan Schmid. Dynamic balanced graph partitioning. *SIAM Journal on Discrete Mathematics*, 34(3):1791–1812, 2020.
- 4 Chen Avin, Kaushik Mondal, and Stefan Schmid. Demand-aware network designs of bounded degree. *Distributed Computing*, 33(3-4):311–325, 2020.
- 5 Yossi Azar, Chay Machluf, Boaz Patt-Shamir, and Noam Touitou. Competitive Vertex Recoloring. In Mikolaż Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2022.13.
- 6 Yossi Azar, Chay Machluf, Boaz Patt-Shamir, and Noam Touitou. Competitive vertex recoloring: (online disengagement). *Algorithmica*, pages 1–27, 2023.
- 7 Luis Barba, Jean Cardinal, Matias Korman, Stefan Langerman, André Van Renssen, Marcel Roeloffzen, and Sander Verdonschot. Dynamic graph coloring. In *Workshop on algorithms and data structures*, pages 97–108. Springer, 2017.
- 8 Sayan Bhattacharya, Fabrizio Grandoni, Janardhan Kulkarni, Quanquan C Liu, and Shay Solomon. Fully dynamic $(\delta + 1)$ -coloring in $o(1)$ update time. *ACM Transactions on Algorithms (TALG)*, 18(2):1–25, 2022.
- 9 Bartłomiej Bosek, Yann Disser, Andreas Emil Feldmann, Jakub Pawlewicz, and Anna Zych-Pawlewicz. Recoloring Interval Graphs with Limited Recourse Budget. In Susanne Albers, editor, *17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2020)*,

- volume 162 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:23, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SWAT.2020.17.
- 10 VMware Docs. Using affinity rules with vsphere drs, 2024. URL: <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-resource-management/GUID-FF28F29C-8B67-4EFF-A2EF-63B3537E6934.html#:~:text=An%20anti%2Daffinity%20rule%20specifies,using%20this%20type%20of%20rule.>
 - 11 Tobias Forner, Harald Räcke, and Stefan Schmid. Online balanced repartitioning of dynamic communication patterns in polynomial time. In *Symposium on Algorithmic Principles of Computer Systems (APOCS)*, pages 40–54. SIAM, 2021.
 - 12 Michael R Garey, David S Johnson, and Larry Stockmeyer. Some simplified np-complete problems. In *Proceedings of the sixth annual ACM symposium on Theory of computing*, pages 47–63, 1974.
 - 13 Monika Henzinger, Stefan Neumann, Harald Räcke, and Stefan Schmid. Tight bounds for online graph partitioning. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2799–2818. SIAM, 2021.
 - 14 Monika Henzinger, Stefan Neumann, and Andreas Wiese. Explicit and implicit dynamic coloring of graphs with bounded arboricity. *arXiv preprint arXiv:2002.10142*, 2020.
 - 15 Monika Henzinger and Pan Peng. Constant-time dynamic $(\delta + 1)$ -coloring. *ACM Trans. Algorithms*, 18(2), March 2022. doi:10.1145/3501403.
 - 16 Fabien Hermenier and Ludovic Henrio. Trustable virtual machine scheduling in a cloud. In *Proceedings of the 2017 Symposium on Cloud Computing*, pages 15–26, 2017.
 - 17 Richard M Karp. *Reducibility among combinatorial problems*. Springer, 2010.
 - 18 Manas Jyoti Kashyop, N. S. Narayanaswamy, Meghana Nasre, and Sai Mohith Potluri. Trade-offs in dynamic coloring for bipartite and general graphs. *Algorithmica*, 85(4):854–878, October 2022. doi:10.1007/s00453-022-01050-7.
 - 19 Henry A Kierstead, Alexandr V Kostochka, Marcelo Mydlarz, and Endre Szemerédi. A fast algorithm for equitable coloring. *Combinatorica*, 30(2):217–224, 2010.
 - 20 Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
 - 21 Maciej Pacut, Mahmoud Parham, and Stefan Schmid. Optimal online balanced graph partitioning. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2021.
 - 22 Harald Räcke, Stefan Schmid, and Ruslan Zabrodin. Approximate dynamic balanced graph partitioning. In *Proceedings of the 34th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 401–409, 2022.
 - 23 Rajmohan Rajaraman and Omer Wasim. Improved bounds for online balanced graph repartitioning. In *30th Annual European Symposium on Algorithms (ESA 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
 - 24 Manoel C Silva Filho, Claudio C Monteiro, Pedro RM Inácio, and Mário M Freire. Approaches for optimizing virtual machine placement and migration in cloud environments: A survey. *Journal of Parallel and Distributed Computing*, 111:222–250, 2018.
 - 25 Shay Solomon and Nicole Wein. Improved dynamic graph coloring. *ACM Trans. Algorithms*, 16(3), June 2020. doi:10.1145/3392724.
 - 26 Ruoxiao Tang. Innovations for data center network optimization. In *2021 International Conference on Computer Information Science and Artificial Intelligence (CISAI)*, pages 610–615, 2021. doi:10.1109/CISAI54367.2021.00123.
 - 27 Huandong Wang, Yong Li, Ying Zhang, and Depeng Jin. Virtual machine migration planning in software-defined networks. *IEEE Transactions on Cloud Computing*, 7(4):1168–1182, 2019. doi:10.1109/TCC.2017.2710193.